IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

In re Application of:	§	Group Art Unit: 2452
	§	
Sreeram Duvvuru	§	Examiner: Chankong, Dohm
	§	
	§	Atty. Dkt. No.: 5681-90800
	§	
	§	
Serial No.: 09/896,244	§	
	§	
Filed: June 28, 2001	§	
	§	
For: Differentiated Quality of Service Context Assignment and Propagation	§	
	§	
	§	
	§	
	§	

APPEAL BRIEF

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir/Madam:

Further to the Notice of Appeal filed April 5, 2011, Appellant presents this Appeal Brief. Appellant respectfully requests that the Board of Patent Appeals and Interferences consider this appeal.

I. REAL PARTY IN INTEREST

The subject application is owned by Sun Microsystems, Inc., which is now Oracle America, Inc.

II. RELATED APPEALS AND INTERFERENCES

No other appeals, interferences or judicial proceedings are known which would be related to, directly affect or be directly affected by or have a bearing on the Board's decision in this appeal.

III. STATUS OF CLAIMS

Claims 1-27 are pending in the application and stand finally rejected. The rejection of claims 1-27 is being appealed. A copy of claims 1-27 is included in the Claims Appendix herein below.

IV. STATUS OF AMENDMENTS

No amendments have been submitted subsequent to the final rejection.

V. SUMMARY OF CLAIMED SUBJECT MATTER

Independent claim 1 is directed to computer-implemented method for providing differentiated quality of service in an application server. The method includes a server system receiving a request for service from a client, where the request includes an encoding specifying a current user role and a requested service. (*See, e.g.,* page 2, lines 20-22; page 10, lines 13-16; and FIG. 4, element 410.) The method also includes, in response to receiving the request for service: accessing pre-determined policy data and establishing a quality of service context based on the specified current user role included in the request and based on the policy data. (*See, e.g.* page 2, lines 21-23; page 10, line 26 - page 11, line 3; page 11 lines 20-22; and FIG. 4, elements 420 and 430.) Also in response to receiving the request for service, the method includes propagating the quality of service context with the request in the server system, where propagating the request includes sending data indicating the quality of service context with the request. (*See, e.g.,* page 2, lines 2-3; page 12, lines 22-25, page 13, lines 2-4; and FIG. 4, element 440.)

Independent claim 19 is directed to a non-transitory computer-readable storage medium. (*See, e.g.,* page 2, lines 24-25.) The storage medium includes program instructions that are executable to implement a server system that is configured to receive a request for service from a client, where the request includes an encoding specifying a current user role and a requested service. (*See, e.g.,* page 10, lines 13-16; page 14, lines 15-17 and 24-28; page 15, lines 4-8; and FIG. 4, element 410.) The server system is also configured to, in response to receiving the request for service, access pre-determined policy data, and establish a quality of service context based on the specified current user role included in the request and based on the policy data. (*See, e.g.,* page 2, lines 25-27; page 10 line 26 – page 11, line 3; page 11, lines 20-22; and FIG. 4, elements 420 and 430.) The server system is also configured to, in response to receiving the request for service, propagate data indicating the quality of service context with the request in the server system. (*See, e.g.,* page 2, line 27; page 12, lines 22-25; page 13, lines 2-4; and FIG. 4, element 440.)

Independent claim 14 is directed to a computer system that includes a processor and a memory storing program instructions. (*See, e.g.,* page 13, line 24 – page 14, line 1; and FIG. 5, computer system 500, processor 504, and memory 506.) The processor is operable to execute the program instructions to implement a server system that is configured to receive a request for service from a client, where the request includes an encoding specifying a current user role and a requested service. (*See, e.g.,* page 2, line 28 – page 3, line 2; page 10, lines 13-16; page 14, lines 15-17; and FIG. 4, element 410.) In response to receiving the request for service, the server system is further configured to access pre-determined policy data and establish a quality of service context based on the specified current user role included in the request and based on the policy data. (*See, e.g.,* page 3, lines 2-3; page 10, line 26 – page 11, line 3; page 11, lines 20-22; and FIG. 4, elements 420 and 430.) In response to receiving the request for service, the server system is further configured to propagate data indicating the quality of service context with the request in the server system. (*See, e.g.,* page. 3, line 4; page 12, lines 22-25; page 13, lines 2-4; and FIG. 4, element 440.)

The summary above describes various examples and embodiments of the claimed subject matter; however, the claims are not necessarily limited to any of these examples and embodiments. The claims should be interpreted based on the wording of the respective claims.

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

1. Claims 1, 2, 4, 5, 7, 9, 10, 11, 13, 14, 16, 18-20, 22, 23, 25 and 27 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Patel et al. (U.S. Patent 6,865,185) (hereinafter "Patel") in view of Win et al. (U.S. Patent 6,453,353) (hereinafter "Win") and further in view of Lupu, et al. ("Use of Roles and Policies for Specifying and Managing a virtual Enterprise") (hereinafter "Lupu").

2. Claims 3, 12 and 21 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Patel, Win and Lupu and further in view of Ayyagari et al. (U.S. Publication 2001/0024434) (hereinafter "Ayyagari").

3. Claims 6, 15 and 24 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Patel, Win and Lupu in view of Zara et al. (U.S. Patent 7,206,848) (hereinafter "Zara").

4. Claims 8, 17 and 26 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Patel, Win and Lupu in view of Vange (U.S. Publication 2002/0059170) (hereinafter "Vange").

VII. ARGUMENT

I. First Ground of Rejection:

The Office Action rejected claims 1, 2, 4, 5, 7, 9, 10, 11, 13, 14, 16, 18-20, 22, 23, 25 and 27 under 35 U.S.C. § 103(a) as being unpatentable over Patel in view of Win and further in view of Lupu. Appellant traverses this rejection for at least the following reasons.

Independent Claim 1:

1. **The cited references clearly fail to teach or suggest a server system receiving a request for service from a client, wherein said request includes an encoding specifying a current user role and a requested service, and in response to receiving the request for service ... establishing a quality of service context based on the specified current user role included in said request and based on said policy data**.

In rejecting claim 1, the Examiner has admitted that Patel does not disclose a server system receiving a request that includes a current user role, and relies on Win to teach this aspect of Appellant's claim. The Examiner has further admitted that Patel does not disclose establishing a quality of service context based on the current user role, and relies on Win and Lupu to teach this aspect of Appellant's claim.

Patel is directed to a system and method for queuing traffic in a wireless communication network. In Patel, packets for transmission include a flow identifier. The packets are assigned to one of a plurality of virtual groups that include discrete transmission resources based on a variety of parameters, including the flow identifier.

Appellant asserts that systems and methods for queuing transmission packets in a communications network are not analogous to (or even in the same filed of endeavor as) systems and methods in which an application server receives and service requests for service from clients or users. On page 2 of the instant Final Action, the Examiner submits that Patel teaches receiving a request for service from a client and that the request includes an encoding specifying the requested service. Specifically, the Examiner submits that Patel discloses packets that include an encoding for a requested service, citing a "service option label" that specifies whether the information being transmitted is from a voice or data call. Appellant notes, however, that this power/service option label 88 is not included in a service request received from a client/user, as the Examiner has suggested, but is inserted into a received transmission packet by the dynamic flow manager 32 based on flow identity. Specifically, it is added to the transmission packet based on information maintained in a call information database (54) about flows in the wireless network. For at least these reasons, Appellant asserts that the transmission packets of Patel do not teach or suggest the service requests of Appellant's claim.

Win is directed to role-based navigation of information resources (e.g., information resources stored on a protected Web server). In Win, user roles (described in Win as "job functions") are used in determining what resources a user can access, not in establishing a quality of service context for access requests or for any other type of service request. The Examiner relies on Win to disclose a request that includes a current user role. Specifically, the Examiner cites column 6, lines 44-48 and 58-65, as teaching, "user sending a request with a cookie identifying the user roles to the server." This and other passages of Win describe that when a user logs into the system (e.g., by logging into an access server using a single sign-on), the user is first authenticated (e.g., using the user's name and password), and then another module of the access server reads the user's roles (plural) from a Registry Server, encrypts them, and sends this information in a "roles cookie" to the user's browser. In other words, in Win, a roles cookie is created by the access server in response to a user login, based on information about the user's roles that is stored in the Registry Server, and then this roles cookie is sent to the user's

browser. This roles cookie is not an encoding specifying one or more current user roles, but is described as containing "a list of the user's roles" (see, e.g., column 10, line 55). There is nothing in Win that describes or implies that the roles cookie specifies which of the user's roles is a current user role (i.e. a role in which the user is currently acting). **Since the roles cookie taught by Win does not specify a current user role, but contains a list of all of the roles that the user is allowed to assume, it cannot be used to establish a QoS context based on the specified current user role.**

The Examiner submits that it would have been obvious to one of ordinary skill in the art to have modified Patel's service requests to include a current user role as taught by Win, stating, "Such a modification is an example of simple substitution of one known element (Win's user request that contains a role cookie) for another (Patel's user request) to obtain predictable results (Patel's system modified to directly receive user roles to identify which policies to apply to the request, see Win, column 5 <<lines 44-54>>)." The Examiner's suggested "predictable result" appears to be that extraneous (and unnecessary) information would be included in the packets transmitted by Patel, since Patel's system does not rely on user roles for any of its functionality, much less for establishing a quality of service context. Instead, Patel uses completely different types of information to determine quality of service parameters for received transmission packets.

The Examiner has admitted that Patel and Win describe role-based access to resources, but do not expressly disclose establishing a quality of service context based on the current user role. The Examiner relies on Lupu to teach this aspect of Appellant's claim. Specifically, the Examiner submits that Lupu discloses the use of the user role to establish particular QoS restraints (e.g., requirements, capabilities, contracts in terms of error rates, throughputs, delay, etc.) on the user's action, citing Lupu, page 1, paragraph 1: Introduction, and page 2, paragraph 2.1: ODP Definition of a role: "A role type… may include additional constraints on the behavior, such as policy or Quality of Service (QoS) statements". **Lupu is directed to the use of roles and policies (e.g., rights and duties associated with an organizational position) for specifying and managing a virtual enterprise, not for receiving requests for service from clients or users, and has**

nothing to do with the limitations of Appellant's claim. For example, the Examiner's citation on page 2 of Lupu does not describe the handling of a service request in a virtual enterprise, much less the use of a user role in establishing a QoS context for a request for service made in (or received by) such a virtual enterprise from a client/user. In fact, Lupu does not include any mention of service requests at all. Instead, the cited passage is actually a description of some of the concepts of the ODP enterprise language that can be used to define and model behavior in a virtual enterprise. In the ODP enterprise language, a "role" is a foundation concept that can be used to define the rights, responsibilities, and/or behavior of various components (objects) of a virtual enterprise. The objects are agents (actors) in the system (human or otherwise). As noted above, a "role type" definition may include QoS statements. However, nothing in Lupu describes receiving or establishing a QoS context for a request for service in the virtual enterprise.

Since Lupu is not directed to establishing QoS context for a request for service, but for defining and using role types to model the behavior of agents in a virtual enterprise, it is not clear how the Examiner means to combine the references in teaching Appellant's claim. For example, the roles cookie taught by Win does not specify a current user role, but contains a list of user roles that a user can assume. Therefore, even if Lupu taught the use of a current user role in establishing a QoS for a request for service (which, since it has nothing to do with service requests or a QoS for handling those requests, it does not) the roles cookie of Win would not specify a current user role (or role type) with which to establish such a QoS context. Instead, each of the user roles listed in the roles cookie of Win would (according to Lupu) be defined using a different role type definition, and could include different QoS statements. In other words, the roles cookie of Win does not specify which of a list of user roles for a given user is a current role, and thus could not be used to map a current user role to a role type definition that includes QoS statements. Therefore, even if there were some way to combine the teachings of Patel, Win, and Lupu, their combination would not result in Appellant's claimed invention.

The Examiner submits, "It would have been obvious to one of ordinary skill in the art to have modified Patel's QoS system to include the user role functionality described above from Win and Lupu. Such a modification would have provided improvement to Patel's system because incorporating role-based QoS (as taught in Win and Lupu) provides a more flexible and extensible way to control QoS in a network [see for example Win, abstract | column 2 <<lines 26-28>>: the user role allows flexibility and extensibility in adding users to the system]." The Examiner's own citation and remarks do not support his conclusions. First, as noted above, Patel is not directed to the receiving (or handling) of client/user service requests, and does not consider user roles in any of its functionality in handling transmission packets for queuing wireless communication traffic. In addition, Win's abstract does not describe an advantage of role-base QoS, as the Examiner suggests, but states, "The registry server controls a flexible, extensible, additive data model stored in a database that describes the user, the resources, roles of the user, and functional groups in the enterprise that are associated with the user." In other words, Win's abstract describes the use (and advantage to the system of Win) of a flexible, extensible, and additive data model. The Examiner's citation in Win, column 2 states, "There is a further need for such a mechanism that is easy to configure and re-configure as new users and resources become part of the system", again referring to the additive data model of Win. The Examiner does not explain how he believes such a data model would improve the system of Patel, nor how the use of role-based QoS would "provide a more flexible and extensible way to control QoS" in the system of Patel, which does not rely on user roles in any of its functions. Appellant asserts that the stated reason to combine the references is not commensurate with the features of Patel, and that the suggested combination would not result in Appellant's claimed invention.

2. When the claim is considered as a whole, the cited art clearly fails to teach or suggest a server system receiving a request for service from a client, wherein said request includes an encoding specifying a current user role and a requested service, and in response to receiving the request for service ... establishing

a quality of service context based on the specified current user role included in said request and based on said policy data.

The Examiner has failed to consider the combination of these limitations in the claim **as a whole**. The Examiner merely submits, in effect, that Patel teaches establishing and propagating a QoS context with a transmission packet, that Win teaches including a list of user roles that a user can assume in a request for service, and Lupu teaches that QoS statements can be included in a role type definition. However, these features are not analogous to the elements of Appellant's claim that the Examiner submits they teach, and they are not combinable to teach or teach or suggest the specific features and functionality in the specific manner combined in Appellant's claim. As noted above, Patel does not teach any use of a user role in determining a QoS context, or in performing any other functions. Therefore, there would be no reason to include such information in its transmission packets. In addition, the roles cookie of Win does not teach or suggest an encoding specifying a current user role. Therefore, even if it were included in a packet in Patel, Patel could not use this information (a list of user roles) to establish a QoS context for such a packet based on the current user role. In addition, there is nothing in Lupu that describes the use of a role type definition (or QoS statements thereof) to establish a QoS context for a request for service received from a client/user, such as the request for service recited in Appellant's claim. Therefore, even if the features of Patel, Win, and Lupu could somehow be combined, they would not result in a system in which a server system receives a request for service from a client, and the request includes an encoding specifying a current user role and a requested service, and in response to receiving the request for service... establishing a quality of service context based on the specified current user role included in said request and based on said policy data. The combined art does not teach the combination of these limitations arranged as recited in claim 1 when the claim is considered in its entirety.

For at least the reasons stated above, Appellant asserts that the Examiner has failed to establish a *prima facie* rejection of claim 1.

Independent Claim 10:

1. The cited references clearly fail to teach or suggest independent claim 10.

Independent claim 10 includes limitations similar to claim 1, and was rejected for the same reasons as claim 1. For example, claim 10 includes the following: *a server system, configured to: receive a request for service from a client, wherein said request includes an encoding specifying a current user role and a requested service; and in response to receiving the request for service... establish a quality of service context based on the specified current user role included in said request and based on said policy data.*

Therefore, the arguments presented above apply with equal force to this claim, as well.

For at least the reasons stated above, Appellant asserts that the Examiner has failed to establish a *prima facie* rejection. The rejection of claim 10 is unsupported by the evidence of record and reversal thereof is respectfully requested.

Independent Claim 19:

1. The cited references clearly fail to teach or suggest independent claim 10.

Independent claim 19 includes limitations similar to claim 1, and was rejected for the same reasons as claim 1. For example, claim 19 includes the following: *a server system configured to: receive a request for service from a client, wherein said request includes an encoding specifying a current user role and a requested service; and in response to receiving the request for service, the server system is further configured to...*

establish a quality of service context based on the specified current user role included in said request and based on said policy data.

Therefore, the arguments presented above apply with equal force to this claim, as well.

For at least the reasons stated above, Appellant asserts that the Examiner has failed to establish a *prima facie* rejection. The rejection of claim 19 is unsupported by the evidence of record and reversal thereof is respectfully requested.

Dependent Claims 2, 4, 5, 7, 9, 11, 13, 14, 16, 18, 20, 22, 23, 25, and 27:

Appellant traverses the rejection of these claims for at least the reasons presented above regarding the independent claims from which they depend.

II. Second Ground of Rejection

The Office Action rejected claims 3, 12 and 21 as being unpatentable over Patel, Win and Lupu and further in view of Ayyagari. Appellant traverses this rejection for at least the following reasons.

Dependent Claim 3:

1. The cited references clearly fail to teach or suggest that said quality of service context includes information indicating service class and a deadline.

In rejecting claim 3, the Examiner submits that Patel discloses a quality of service context that includes a deadline. Specifically, the Examiner cites paragraph [0006] as disclosing, "execution of a desired task in a specified time period." The cited passage

actually states, "QoS refers to a reservation of resources, such as bandwidth, time slices, relative priority, memory, ports and the like that are required for the execution of a desired task in a specified time period. Default QoS level, typically termed 'best effort,' represents execution of a task if resources are available when needed, but not necessarily providing reservation of resources. In other words, 'best effort' represents providing otherwise idle resources for carrying out the task. Higher levels of assurance provide better than 'best effort' and can include several levels of relative priority as is discussed next." In other words, the QoS level for a task refers to whether resources needed to carry out a task in a specified time period would be made available for that task. The QoS level itself does not specify the deadline for carrying out a task. The Examiner also cites paragraph [0048] as disclosing a "time constraint". This passage states, "The laptop computer 215, which is one of the wireless devices 210, 215 or 220, requests access to a network to communicate, with a specified QoS specifying the required bandwidth, time constraints and the like, with a receiving node 230. In response, the access point 200 sends a message requesting QoS to a router 235. The router 235 manages packet flow through a subnet 240 to allow the requested communication access to the receiving node 230." This general reference to "time constraints" does not explicitly teach that a specific "deadline" would be specified by the QoS level for a given transmission packet.

For at least the reasons stated above, Appellant asserts that the Examiner has failed to establish a *prima facie* rejection. The rejection of claim 3 is unsupported by the evidence of record and reversal thereof is respectfully requested.

Dependent Claim 12:

Dependent claim 12 includes limitations similar to those of claim 3, and was rejected for the same reasons as claim 3. For example, claim 12 includes the following: *said quality of service context includes information indicating service class and a deadline.*

Therefore, the arguments presented above and directed to claim 3 apply with equal force to this claim, as well.

For at least the reasons stated above, Appellant asserts that the Examiner has failed to establish a *prima facie* rejection. The rejection of claim 12 is unsupported by the evidence of record and reversal thereof is respectfully requested.

Dependent Claim 21:

Dependent claim 21 includes limitations similar to those of claim 3, and was rejected for the same reasons as claim 3. For example, claim 21 includes the following: *said quality of service context includes information indicating service class and a deadline*.

Therefore, the arguments presented above and directed to claim 3 apply with equal force to this claim, as well.

For at least the reasons stated above, Appellant asserts that the Examiner has failed to establish a *prima facie* rejection. The rejection of claim 21 is unsupported by the evidence of record and reversal thereof is respectfully requested.

III. Third Ground of Rejection

The Office Action rejected claims 6, 15 and 24 as being unpatentable over Patel, Win and Lupu in view of Zara. Appellant traverses this rejection for at least the following reasons.

Dependent Claim 6:

1. The cited references clearly fail to teach or suggest propagating the same quality of service context with a subsequent sub-request of said request.

The Examiner has admitted that Patel as modified by Win and Lupu does not expressly disclose propagating the same quality of service content with a subsequent request, and relies on Zara to teach this aspect of Appellant's invention. **Appellant's first note that the Examiner has improperly ignored the term "sub-request" in the claim.** "All words in a claim must be considered in judging the patentability of that claim against the prior art." MPEP 2143.03; *In re Wilson*, 424 F.2d 1382, 1385, 165 USPQ 494, 496 (CCPA 1970).

The Examiner cites column 7, lines 58-61 of Zara as disclosing attaching the same quality of service context ("tag") with a subsequent request related to the first request. Zara is directed to intelligent classification and handling of user requests in a data service system. The cited passage (and preceding text) states that when a client system generates a first request for a session or transaction, the first request does not contain a classification tag. Instead, a classification tag is generated by the application system (according to business rules stored in the application system) and is attached to the response to the first request and sent back to the client that sent the first request. The requesting client attaches the classification tag that was generated and sent with the first response to any subsequent requests for the same session or transaction. When the application system receives the subsequent requests, it reapplies the business rules and may re-classify them and attach a different classification tag to the corresponding responses (see, e.g., FIG. 3 and its description in columns 8-9).

Appellant submits that, as described above, Zara does not describe propagating the same quality of service context with a subsequent sub-request of (an original) request. In fact, Zara does not describe the propagation of classification information from one request to another request at all, but from one response to a next request. In addition, Zara teaches that for each subsequent request, the business rules are re-evaluated, which may result in a different classification for each subsequent request. Therefore, Zara clearly does not teach or suggest the above-referenced limitations of Appellant's claim.

The Examiner submits that it would have been obvious to one of ordinary skill in the art to have modified Patel to include Zara's teachings to insure that the requests involved in the same session or transaction receive the QoS. However, as noted above, Zara does not teach that requests involved in the same session or transaction receive the same classification. Instead, in Zara: no classification is included in a first request; classifications do not propagate from one request to another request, but from one response to a next request; and since business rules are re-evaluated for each request, there is nothing to ensure that requests involved in the same session or transaction receive the same classification. Therefore, the Examiner's stated motivation to combine the references is inconsistent with the teachings of the references themselves.

Appellant notes that the Examiner did not address the arguments above in the Response to Arguments section of the instant Final Action.

For at least the reasons stated above, Appellant asserts that the Examiner has failed to establish a *prima facie* rejection of claim 6.

Dependent Claim 15:

Dependent claim 15 includes limitations similar to those of claim 6, and was rejected for the same reasons as claim 6. For example, claim 15 includes the following: *program instructions executable to: propagate the same quality of service context with a subsequent sub-request of said request.*

Therefore, the arguments presented above and directed to claim 6 apply with equal force to this claim, as well.

For at least the reasons stated above, Appellant asserts that the Examiner has failed to establish a *prima facie* rejection. The rejection of claim 15 is unsupported by the evidence of record and reversal thereof is respectfully requested.

Dependent Claim 24:

Dependent claim 24 includes limitations similar to those of claim 6, and was rejected for the same reasons as claim 6. For example, claim 24 includes the following: *program instructions to: propagate the same quality of service context with a subsequent sub-request of said request*.

Therefore, the arguments presented above and directed to claim 6 apply with equal force to this claim, as well.

For at least the reasons stated above, Appellant asserts that the Examiner has failed to establish a *prima facie* rejection. The rejection of claim 24 is unsupported by the evidence of record and reversal thereof is respectfully requested.

IV. Fourth Ground of Rejection

The Office Action rejected claims 8, 17 and 26 as being unpatentable over Patel, Win and Lupu in view of Vange. Appellant traverses this rejection for at least the following reasons.

Dependent Claim 8:

1. The cited art clearly fails to teach or suggest *wherein said propagating comprises a load balancing service dispatching said request, including said quality of service context, to an application server in a plurality of application servers in the server system, based on said quality of service context.*

In rejecting claim 8, the Examiner has admitted that Patel does not expressly disclose a load balancing service that dispatches requests to an application server and relies on Vange to teach this aspect of Appellant's invention. Vange is directed to a system for load balancing between web servers in a network environment. The Examiner

submits that Vange discloses a load balancing service that dispatches requests to an application server in a plurality of application servers based on a quality of service context, citing FIG. 1 of Patel, FIG. 2 of Vange, and paragraph [0094] of Vange. Appellant asserts however, that the QoS data used in load balancing in Vange is not a QoS context included in a request for service, but is data about current conditions collected using other means, such as data collected by QoS monitor 64 regarding the status of the system with respect to the current workload.

The Examiner submits that it would have been obvious to one of ordinary skill in the art to have modified Patel to include Vange's load balancing capability to insure that loads are balanced equally between the servers. Appellant first asserts that Patel includes other mechanisms for metering, adaptive congestion control, and flow control that facilitate a fair and equitable distribution of packets to virtual groups and to resources allocated to those virtual groups that take into account available bandwidth, load, and QoS class (see, e.g., FIG, 15, and its description). Therefore, there would be no reason to look to Vange for another method of balancing packet distribution, nor is it clear that the load balancing of Vange would necessarily improve Patel's capability to distribute packets to virtual groups and their resources. In addition, since Vange's load balancing system does not rely on a QoS context of a particular service request, nor does it propagate such a QoS context in a service request to an application server, even if the load balancing taught by Vange were incorporated into the system of Patel, the combination would not result in Appellant's claimed invention. Instead, it would merely add a load balancing function based on monitored performance data to the system of Patel.

Appellant notes that the Examiner did not address the arguments above in the Response to Arguments section of the instant Final Action.

For at least the reasons stated above, Appellant asserts that the Examiner has failed to establish a *prima facie* rejection of claim 8.

Dependent Claim 17:

Dependent claim 17 includes limitations similar to those of claim 8, and was rejected for the same reasons as claim 8. For example, claim 17 includes the following: *wherein said propagating comprises a load balancing service dispatching said request, including said quality of service context, to an application server in a plurality of application servers in the server system, based on said quality of service context.*

Therefore, the arguments presented above and directed to claim 8 apply with equal force to this claim, as well.

For at least the reasons stated above, Appellant asserts that the Examiner has failed to establish a *prima facie* rejection. The rejection of claim 17 is unsupported by the evidence of record and reversal thereof is respectfully requested.

Dependent Claim 26:

Dependent claim 26 includes limitations similar to those of claim 8, and was rejected for the same reasons as claim 8. For example, claim 26 includes the following: *wherein said propagating comprises a load balancing service dispatching said request including, said quality of service context, to an application server in a plurality of application servers in the server system, based on said quality of service context.*

Therefore, the arguments presented above and directed to claim 8 apply with equal force to this claim, as well.

For at least the reasons stated above, Appellant asserts that the Examiner has failed to establish a *prima facie* rejection. The rejection of claim 26 is unsupported by the evidence of record and reversal thereof is respectfully requested.

CONCLUSION

For the foregoing reasons, it is submitted that the Examiner's rejection of claims 1-27 was erroneous, and reversal of the rejections is respectfully requested.

The Commissioner is authorized to charge the appeal brief fee and any other fees that may be due to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5681-90800/RCK.

Respectfully submitted,

/Robert C. Kowert/
Robert C. Kowert, Reg. #39,255
Attorney for Appellant

Meyertons, Hood, Kivlin, Kowert & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
(512) 853-8850

Date: June 6, 2011

VIII. CLAIMS APPENDIX

The claims on appeal are as follows.

1. A computer-implemented method for providing differentiated quality of service in an application server, comprising:

a server system receiving a request for service from a client, wherein said request includes an encoding specifying a current user role and a requested service; and

in response to receiving the request for service:

accessing pre-determined policy data;

establishing a quality of service context based on the specified current user role included in said request and based on said policy data; and

propagating said quality of service context with said request in the server system, wherein said propagating comprises sending data indicating the quality of service context with the request.

2. The method of claim 1, wherein said request further includes information indicating at least one of a time constraint or a user identity.

3. The method of claim 1 wherein said quality of service context includes information indicating service class and a deadline.

4. The method of claim 1 wherein said establishing a quality of service context is completed at an ingress point.

5. The method of claim 4 wherein said ingress point is at least one of a web server or a protocol manager service within said server system.

6. The method of claim 1 further comprising, propagating the same quality of service context with a subsequent sub-request of said request.

7. The method of claim 1 wherein said propagating includes inserting said quality of service context into data sent with the request adjacent to at least one of a security and transaction context.

8. The method of claim 1, wherein said propagating comprises a load balancing service dispatching said request, including said quality of service context, to an application server in a plurality of application servers in the server system, based on said quality of service context.

9. The method of claim 1, wherein said propagating comprises a request manager service dispatching said request, including said quality of service context, to a software component in a plurality of software components in the server system, based on said quality of service context.

10. A non-transitory computer-readable storage medium, comprising program instructions executable to implement:

a server system, configured to:

receive a request for service from a client, wherein said request includes an encoding specifying a current user role and a requested service; and

in response to receiving the request for service:

access pre-determined policy data;

establish a quality of service context based on the specified current user role included in said request and based on said policy data; and

propagate data indicating said quality of service context with said request in the server system.

11. The non-transitory computer-readable storage medium of claim 10, wherein said request further includes information indicating at least one of a time constraint or a user identity.

12. The non-transitory computer-readable storage medium of claim 10, wherein said quality of service context includes information indicating service class and a deadline.

13. The non-transitory computer-readable storage medium of claim 10, wherein said establishing a quality of service context is completed at an ingress point.

14. The non-transitory computer-readable storage medium of claim 13, wherein said ingress point is at least one of a web server or a protocol manager service within said server system.

15. The non-transitory computer-readable storage medium of claim 10, further comprising program instructions executable to: propagate the same quality of service context with a subsequent sub-request of said request.

16. The non-transitory computer-readable storage medium of claim 10, wherein said propagating includes inserting said quality of service context into data sent with the request adjacent to at least one of a security and transaction context.

17. The non-transitory computer-readable storage medium of claim 10, wherein said propagating comprises a load balancing service dispatching said request, including said quality of service context, to an application server in a plurality of application servers in the server system, based on said quality of service context.

18. The non-transitory computer-readable storage medium of claim 10, wherein said propagating comprises a request manager service dispatching said request, including said quality of service context, to a software component in a plurality of software components in the server system, based on said quality of service context.

19. A first computer system, comprising:

a processor;

a memory storing program instructions;

wherein the processor is operable to execute the program instructions to implement a server system configured to:

receive a request for service from a client, wherein said request includes an encoding specifying a current user role and a requested service; and

in response to receiving the request for service, the server system is further configured to:

access pre-determined policy data;

establish a quality of service context based on the specified current user role included in said request and based on said policy data; and

propagate data indicating said quality of service context with said request in the server system.

20. The system of claim 19, wherein said request further includes information indicating at least one of a time constraint or a user identity.

21. The system of claim 19, wherein said quality of service context includes information indicating service class and a deadline.

22. The system of claim 19, wherein said establishing a quality of service context is completed at an ingress point.

23. The system of claim 22, wherein said ingress point is at least one of a web server or a protocol manager service within said server system.

24. The system of claim 19, further comprising program instructions to: propagate the same quality of service context with a subsequent sub-request of said request.

25. The system of claim 19, wherein said propagating includes inserting said quality of service context into data sent with the request adjacent to at least one of a security and transaction context.

26. The system of claim 19, wherein said propagating comprises a load balancing service dispatching said request including, said quality of service context, to an

application server in a plurality of application servers in the server system, based on said quality of service context.

27. The system of claim 19, wherein said propagating comprises a request manager service dispatching said request, including said quality of service context, to a software component in a plurality of software components in the server system, based on said quality of service context.

IX. EVIDENCE APPENDIX

No evidence submitted under 37 CFR §§ 1.130, 1.131 or 1.132 or otherwise entered by the Examiner is relied upon in this appeal.

X. RELATED PROCEEDINGS APPENDIX

There are no related proceedings.